

VMware Brain Deployment Guide

Version: October 7, 2025

Table of Contents

Introduction	2
About VMware Brain Images.....	2
VMware Brain Requirements and Throughput.....	2
Connectivity Requirements (Firewall Rules)	4
Vectra Cloud Connectivity	4
Connectivity Requirements – General	8
Licensing and Deployment Overview.....	9
Licensing Enforcement.....	9
Deployment Overview	9
Brain Deployment in VMware.....	10
Requirements	10
Downloading the latest VMware Brain OVA image	11
Deploying the OVA (vSphere Client / vCenter Server).....	11
Deploying the OVA (embedded host client for ESXi).....	12
32 Core NUMA Configuration	13
Initial boot up and licensing	14
Post Deployment Guidance	18
Setting a static IP and DNS after initial DHCP deployment	18
Required Sensor update package for offline mode (Quadrant UX only)	19
Performance testing.....	20
Integrity checks	20
Configuration validation	21
License checks and renewal.....	21
About VMware vCenter Integration.....	23
Enabling vCenter Integration.....	25
Resizing the Brain	26
Next Steps.....	26
Worldwide Support Contact Information	26

Introduction

This guide is intended to help customers or partners deploy a virtual Brain appliance in VMware environments. A VMware Brain appliance can be used in Vectra AI Platform deployments that use either the Respond UX or the Quadrant UX. The Respond UX is served from Vectra’s cloud and the Quadrant UX is served locally from the Brain appliance. For more detail on Respond UX vs Quadrant UX please see [Vectra Analyst User Experiences \(Respond vs Quadrant\)](#).

This guide will cover basic background information, connectivity requirements (firewall rules that may be needed in your environment), licensing, deployment, and next steps. One of the below guides should be the starting point for your overall Vectra deployment:

- ▼ [Vectra Respond UX Deployment Guide](#)
- ▼ [Vectra Quadrant UX Deployment Guide](#)

Either of the above guides cover basic firewall rules needed for the deployment and initial platform settings. Virtual Sensor (VMware, Hyper-V, KVM, AWS, Amazon, and GCP) configuration and pairing and covered in [their respective guides](#). Physical appliance pairing is covered in the [Vectra Physical Appliance Pairing Guide](#). Please see the [Vectra Product Documentation Index](#) on the Vectra support site for additional documentation including deployment guides for [CDR for M365 / IDR for Azure AD](#) and [CDR for AWS](#).

About VMware Brain Images

The .ova image used to deploy a Brain in VMware is made available on the [Vectra Customer Portal](#) which is part of [Vectra Support](#). Vectra periodically updates the base image used for VMware Brain deployment. It is a best practice to always download the latest image from the Vectra Customer Portal prior to deployment of a new VMware Brain. Brains that are connected to Vectra are updated automatically according to the settings on that Brain. Offline updates are also possible for Quadrant UX deployments only. Please see the [Vectra Quadrant UX Deployment Guide](#) for more details regarding offline updates. Please see the [Offline Updates](#) article on the Vectra Support site for instructions on how to apply offline updates.

VMware Brain Requirements and Throughput

For use in Respond UX or Quadrant UX deployments:

Resource Type	Requirement		
Performance ¹	2 Gbps	4 Gbps	10 Gbps
CPU	8 Cores	16 Cores	32 Cores ³
Memory	64 GB RAM	128 GB RAM	256 GB RAM
Drive (OS, Data) Requires 260 MB/s	128 GB, 512 GB	128 GB, 512 GB	128 GB, 512 GB
Max Paired Sensors	15	25	100
Max Simultaneous Tracked Hosts ²	50,000	50,000	150,000
vSphere Supported Versions	6.5 through 8 *		

For use ONLY in Respond UX for Network deployments (using network Sensors with the Respond UX):

Resource Type	Requirement (Respond UX for Network Deployments Only)	
Performance ¹	150 Mbps	500 Mbps
CPU	4 Cores	6 Cores
Memory	48 GB RAM	48 GB RAM
Drive (OS, Data) Requires 260 MB/s	128 GB, 512 GB	128 GB, 512 GB
Max Paired Sensors	5	10
Max Simultaneous Tracked Hosts ²	25,000	37,500
vSphere Supported Versions	6.5 through 8 *	

¹ Performance represents the aggregate bandwidth observed on the capture interfaces of any Sensors that are paired to the Brain. Guidance is for average traffic mixes. Traffic mixes that skew toward larger flows (like file transfers) will perform better than traffic mixes that skew towards smaller flows (like DNS) as they produce more metadata.

² Refers to how many hosts the Brain can track simultaneously (open host sessions). Brains retain and display data for larger numbers of hosts, this only refers to how many hosts the system can process metadata for simultaneously.

³ Please see [32 Core NUMA Configuration](#) for details on checking and setting (if required) for 32 core vSensors

▼ *** Special Note: Vectra supported VMware hardware versions**

- Vectra supports only versions 11 and 15 of VMware hardware.
- **DO NOT** update the hardware version if offered during deployment or any other situation.
- If you move to an unsupported hardware version, contact Vectra support for guidance. Downgrades may be possible but are not officially supported. Support will be best effort in these situations.
- ▼ The **virtual CPU MUST** support the pdpe1gb cpu flag (1GB Large Pages) – [More information](#), and a minimum SSE instruction level of 4.2, and must support the POPCNT (population count) instruction. This requires the hypervisor host to be running one of the following processors or later:
 - Intel Nehalem (2008) processors and newer
 - AMD Bulldozer (2011) processors and newer
 - Check [VMware's Enhanced vMotion Compatibility \(EVC Explained\) article](#) for details on EVC settings that may mask the underlying physical CPU's required flags. Change EVC settings if required.
- ▼ Vectra VMware based Brains do not support Mixed Mode deployment. They can only be used in Brain mode.
- ▼ Vectra VMware based Brains support running in FIPS mode. Note that the underlying hardware must also be FIPS compliant (it must support the RDRAND CPU instruction).
- ▼ Vectra VMware based Brains and Sensors do not support DirectPath or SR-IOV passthrough.
- ▼ Vectra VMware based Brains and Sensors do not support emulated network adapters.
- ▼ Vectra VMware based Brains and Sensors do support paravirtualized NIC. Vectra uses VMXNET3 ports.
- ▼ Vectra recommends that Brains are configured to use storage local to the hypervisor and are not stored on a SAN. Vectra Brains require extremely high throughput from their disk storage and this throughput cannot normally be sustained by SAN systems without impact to other SAN users.
- ▼ See [Virtual Sensor \(vSensor\) specifications](#) for additional guidance around Storage/SANs, networking requirements, vMotion, Enhanced vMotion compatibility, and unsupported hypervisors. This article was written from the perspective of virtual Sensors, but the considerations also apply to VMware Brains.
- ▼ vMotion is compatible with VMware Brains but new HW or copying the VM can cause VMware to generate a new UUID which causes the Brain license to become invalid, and require relicensing.
 - If VMware gives you a choice, to keep the existing UUID, always pick "I moved it" or "Keep it" instead of copying to retain the UUID and avoid relicensing. See these VMware KBs for more details:
 - [Changing or keeping a UUID for a moved virtual machine](#), [Migrating VMs with vSphere vMotion](#)

Connectivity Requirements (Firewall Rules)

The Vectra AI Platform uses several TCP/UDP ports for different communication purposes. This document will detail basic requirements for initial setup and pairing. Many features and integrations are optional and not in scope for this guide. Additional connectivity requirement guidance is in the [Vectra Respond UX Deployment Guide](#) or [Vectra Quadrant UX Deployment Guide](#). For full detail on all possible firewall rules that might be required in your environment please see [Firewall requirements for Vectra appliances](#) on the Vectra support site.

Vectra Cloud Connectivity

- ▼ For this document, the portions of the Vectra AI Platform that reside in Vectra’s cloud are referred to as the Vectra cloud.
 - This does not refer to any specific service offering.
- ▼ Please check each category below to see if it is applicable to your deployment and if rules are required in your environment to enable the required connectivity.
 - For rule categories that have multiple region options, it is only necessary to put rules in place to allow connectivity to the region that your Vectra tenant is deployed in. This region should be visible in the URL used to access the Respond UX.
 - i.e. [tenant_id].ew1.prod.vectra-svc.ai is used for EU deployments (ew1).
- ▼ RUX for Network refers to a RUX deployment that has enabled network data sources (sensors).
 - This means you have a Brain somewhere in your premises (data center or public cloud) that is paired with network sensors (virtual or physical) to capture network traffic and distill a metadata stream for processing by the Brain appliance.
 - Please refer to the [Vectra Respond UX Deployment Guide](#) for more details.
- ▼ Please refer to the table below to see applicability of the various categories.
- ▼ The “For Brain or User’s Host” column should be interpreted as follows:
 - Brain – Rules required for the Brain to the Vectra Cloud.
 - User’s Web Browser – Rules required for the User’s Web Browser to the Vectra cloud.

Rule Category	Required For	For Brain or User’s Host
RUX for Network GUI Synchronization	RUX for Network Deployments	Brain
Auth Gateways	RUX for Network Deployments Quadrant UX deployments of CDR for M365, IDR for Azure AD, and CDR for AWS.	Brain
RUX Metadata Forwarding	RUX for Network Deployments	Brain
RUX Research Metadata Forwarding	RUX for Network Deployments	Brain
RUX Analyst/Admin Access	All RUX Deployments	User’s Web Browser
RUX Static Asset CDN	All RUX Deployments	User’s Web Browser
RUX Customer File Upload	All RUX Deployments	User’s Web Browser

RUX for Network GUI Synchronization

- ▼ **Required for:** All RUX for Network deployments.
- ▼ This is used to synchronize configurations between the Brain appliance and your Vectra tenant.
- ▼ This communications channel is initiated from the Brain to the endpoint in your Vectra tenant's region.
- ▼ The protocol and ports in use for each entry is the same: **Websocket and HTTPS over TCP/443**

Fully Qualified Domain Name (FQDN)	IP(s)	Region	Initiated From
main-cbi-tunnel-uw2.app.prod.vectra-svc.ai	Dynamic	US	Brain
main-cbi-tunnel-ew1.app.prod.vectra-svc.ai	Dynamic	EU	Brain
main-cbi-tunnel-ec2.app.prod.vectra-svc.ai	Dynamic	Switzerland	Brain
main-cbi-tunnel-cc1.app.prod.vectra-svc.ai	Dynamic	Canada	Brain
main-cbi-tunnel-as2.app.prod.vectra-svc.ai	Dynamic	Australia	Brain

Auth Gateways

- ▼ **Required for:**
 - All Respond UX for Network Deployments.
 - Quadrant UX deployments of CDR for M365, IDR for Azure AD, and CDR for AWS.
 - Your Brain must be able to securely access the Vectra cloud over TCP/443 HTTPS connections to enable detection events from these products to be reported to your UI.
- ▼ In Respond UX for Network deployments, the Brain forwards network detections, entities, host sessions, and any selective PCAPs (Vectra Packet Capture) to your Vectra tenant via this connection.
- ▼ This communications channel is initiated from your Brain to the endpoint in your Vectra tenant's region.

Fully Qualified Domain Name (FQDN)	IP(s)	Protocol / Ports	Region	Initiated From
authgateway.uw2.public.app.prod.vectra-svc.ai	54.245.33.175 52.42.70.176 100.21.109.72 52.26.91.157	HTTPS - TCP/443	US	Brain
authgateway.ew1.public.app.prod.vectra-svc.ai	54.171.40.108 54.246.213.148 54.75.47.147	HTTPS - TCP/443	EU	Brain
authgateway.ec2.public.app.prod.vectra-svc.ai	16.62.18.237 16.62.142.98 51.96.54.201	HTTPS - TCP/443	Switzerland	Brain
authgateway.cc1.public.app.prod.vectra-svc.ai	3.96.112.208 52.60.211.221 15.222.69.161	HTTPS - TCP/443	Canada	Brain
authgateway.as2.public.app.prod.vectra-svc.ai	13.54.11.66 13.55.79.24 13.55.106.102	HTTPS - TCP/443	Australia	Brain

RUX Metadata Forwarding

- ▼ **Required for:** All Respond UX for Network Deployments.
- ▼ Network metadata is forwarded to AWS S3 buckets and processed to make it available for features such as Instant Investigation and Advanced Investigation in the Respond UX.
- ▼ This communications channel is initiated from your Brain to the endpoint in your Vectra tenant's region.
- ▼ The protocol and ports in use for each entry is the same: **HTTPS over TCP/443**

Fully Qualified Domain Name (FQDN)	IP(s)	Region	Initiated From
cbo-upload-network-metadata-forwarder-uswt2-371371611652.s3-accesspoint.us-west-2.amazonaws.com	Dynamic	US	Brain
cbo-upload-network-metadata-forwarder-euwt1-371371611652.s3-accesspoint.eu-west-1.amazonaws.com	Dynamic	EU	Brain
cbo-upload-network-metadata-forwarder-eucl2-371371611652.s3-accesspoint.eu-central-2.amazonaws.com	Dynamic	Switzerland	Brain
cbo-upload-network-metadata-forwarder-cacl1-371371611652.s3-accesspoint.ca-central-1.amazonaws.com	Dynamic	Canada	Brain
cbo-upload-network-metadata-forwarder-apse2-371371611652.s3-accesspoint.ap-southeast-2.amazonaws.com	Dynamic	Australia	Brain

RUX Research Metadata Forwarding

- ▼ **Optional but highly recommended for:** All Respond UX for Network Deployments.
- ▼ Research metadata from precursor algorithms are used to improve model quality and reduce detection noise.
- ▼ This communications channel is initiated from your Brain to the endpoint in your Vectra tenant's region.
- ▼ The protocol and ports in use for each entry is the same: **HTTPS over TCP/443**

Fully Qualified Domain Name (FQDN)	IP(s)	Region	Initiated From
cbo-upload-network-precursors-uswt2-371371611652.s3-accesspoint.us-west-2.amazonaws.com	Dynamic	US	Brain
cbo-upload-network-precursors-euwt1-371371611652.s3-accesspoint.eu-west-1.amazonaws.com	Dynamic	EU	Brain
cbo-upload-network-precursors-eucl2-371371611652.s3-accesspoint.eu-central-2.amazonaws.com	Dynamic	Switzerland	Brain
cbo-upload-network-precursors-cacl1-371371611652.s3-accesspoint.ca-central-1.amazonaws.com	Dynamic	Canada	Brain
cbo-upload-network-precursors-apse2-371371611652.s3-accesspoint.ap-southeast-2.amazonaws.com	Dynamic	Australia	Brain

RUX Analyst/Admin Access

- ▼ **Required for:** All Respond UX deployments.
- ▼ Any analyst or admin that wishes to access the Respond UX will need to ensure that their browser can reach their Vectra tenant to login and access the UI.
- ▼ This communications channel is initiated from the user's host.
- ▼ The protocol and ports in use for each entry is the same: **HTTPS over TCP/443**

Fully Qualified Domain Name (FQDN)	IP(s)	Region	Initiated From
[tenant_id].uw2.portal.vectra.ai	Dynamic	US	User's Web Browser
[tenant_id].ew1.portal.vectra.ai	Dynamic	EU	User's Web Browser
[tenant_id].ec2.portal.vectra.ai	Dynamic	Switzerland	User's Web Browser
[tenant_id].cc1.portal.vectra.ai	Dynamic	Canada	User's Web Browser
[tenant_id].as2.portal.vectra.ai	Dynamic	Australia	User's Web Browser

RUX Static Asset CDN

- ▼ **Required for:** All Respond UX deployments.
- ▼ The Respond UX has certain static assets (HTML, CSS, JS) that are required to serve the web application hosted by a CDN (Content Delivery Network).
- ▼ This communications channel is initiated from the user's host.

Fully Qualified Domain Name (FQDN)	Protocol / Ports	IP(s)	Region	Initiated From
dd6462tdmvp79.cloudfront.net dpew7prsvwb0.cloudfront.net	HTTPS - TCP/443	Dynamic	All	User's Web Browser

RUX Customer File Upload

- ▼ **Required for:** All Respond UX deployments.
- ▼ This communications channel is used for:
 - Vectra Match deployments and will allow upload of rulesets.
 - PCAP download from the Vectra Cloud for Selective PCAP (Vectra Packet Capture)
 - Additional capabilities are planned for future releases.
 - It is recommended to put rules in place even if you don't use Match or Selective PCAP.
- ▼ This communications channel is initiated from the user's host.

Fully Qualified Domain Name (FQDN)	Protocol / Ports	IP(s)	Region	Initiated From
prd-main-customerfiles-580786928539-uswt2.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	US	User's Web Browser
prd-main-customerfiles-580786928539-euwt1.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	EU	User's Web Browser
prd-main-customerfiles-580786928539-eucl2.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	Switzerland	User's Web Browser
prd-main-customerfiles-580786928539-cacl1.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	Canada	User's Web Browser
prd-main-customerfiles-580786928539-apse2.s3.amazonaws.com	HTTPS - TCP/443	Dynamic	Australia	User's Web Browser

Connectivity Requirements – General

Source	Destination	Protocol/Port	Description	QUX-RUX-Both
Admin hosts	Brain / Sensors	TCP/22 (SSH)	CLI access for Brain and Sensors.	Both
Admin hosts	Brain	TCP/443 (HTTPS)	Web UI of Brain appliances (Quadrant UX). Redirect / Status of Brain (Respond UX).	Both
Brain	update2.vectranetworks.com (54.200.156.238)	TCP/443 (HTTPS)	Automatic updates. Pairing keys for physical sensors.	Both
Brain	api.vectranetworks.com (54.200.5.9)	TCP/443 (HTTPS)	Health monitoring, algorithm support, reverse lookups for external IPs, Vectra Threat Intelligence, additional detection content.	Both
Brain	rp.vectranetworks.com (54.200.156.238)	TCP/443 (HTTPS)	Used only for Brains deployed in IaaS clouds. Used for authentication and verification (integrity check of the file system).	Both
Brain	rs.vectranetworks.com (74.201.86.229)	TCP/443 and UDP/9970	Remote Support. OpenVPN type if using firewall with App ID rules.	Both
Brain	DNS servers (as configured)	TCP/53, UDP/53	Both TCP and UDP are required for normal operation.	Both
Brain	NTP servers (as configured) Default is ntp.ubuntu.com	UDP/123	Time synchronization.	Both
Brain	SMTP servers (as configured)	TCP (as configured)	Email alerting (optional but suggested).	Quadrant UX
Sensors, Stream	Brain	TCP/22 (SSH), TCP/443 (HTTPS)	Pairing, metadata transfer, and ongoing communication.	Both
Brain	Sensors, Stream	TCP/22 (SSH)	Remote management and troubleshooting.	Both
Brain	Recall collector	TCP/443 (HTTPS)	Destinations provisioned after enabled.	Quadrant UX
Brain	metadata.vectra.ai (100.20.236.31, 44.229.57.246, 44.228.37.60, 44.228.101.87)	TCP/443 (HTTPS)	Optional anonymized metadata sharing to contribute to future algorithm development.	Quadrant UX

- ▼ Customers should note that the following IP ranges will conflict with remote support capability:
 - 192.168.72.0/21 and 192.168.80.0/21
 - For remote support outside of screen sharing sessions, care should be taken to number the management network interface (MGT) used on any appliance (Brains and Sensors) outside of the above ranges. If your management network interface (MGT) is numbered in either of these ranges, remote support access will not function.

Licensing and Deployment Overview

Appliance code has been encrypted to protect Vectra’s intellectual property and a license is required to enable successful decryption of the file system and deployment. The licensing for Vectra NDR (formerly Detect for Network) running on VMware Brains also governs the ability of the system to create Detections. After deployment, if your Vectra NDR license expires, Detection algorithms will cease operation until a valid license is applied. Recall and Stream operation are unaffected if your Vectra NDR license expires and your Recall and/or Stream license is valid.

Deployment of a Brain appliance in VMware environments can be done on both standalone ESXi servers using the embedded host client or in full vSphere environments using vCenter to manage your virtual infrastructure. If deploying on standalone ESXi servers using their embedded host client, configuring a static IP is not possible until after the deployment is completed and you have access to the Brain CLI. DHCP is the only supported IP assignment method while using the embedded host client. Deployment using vCenter allows for Static or DHCP IP assignment of the Management port for the Brain appliance.

Licensing Enforcement

NDR (Detect for Network) versions 6.20 and higher support new licensing functionality regardless of the type of deployment (physical appliance, cloud IaaS, VMware). All versions will be able to see license status and enable requests for and application of licenses. Enforcement of NDR licensing is only enabled on VMware Brains. Any other Brain type does NOT currently have licensing for NDR enforced. Vectra does plan to add licensing enforcement for other Brain types in the future. It is recommended that all customers work with their account teams to ensure their licensing is up to date. Please refer to the following table for additional detail:

Product	Deployment Type	License Enforcement
NDR (Detect for Network)	VMware Brain	Algorithms stop producing Detections when expired.
NDR (Detect for Network)	Physical or Cloud Brains	Not currently enforced. Planned for future (timing TBD).

Other Vectra products such as Recall, Stream, or CDR for M365 and IDR Azure AD are also licensed but enforcement of the license is a matter of contract compliance between sales teams and customers or partners.

Deployment Overview

The main steps for the deployment are summarized below. For additional detail see [Brain Deployment in VMware](#).

- ▼ Download the .OVA Brain appliance image from <https://support.vectra.ai/vectra/additional-resources>
 - You must be logged into your Vectra support account, to see the download option.
- ▼ Deploy OVA in VMware and power on the appliance using one of the below methods.
 - Using the embedded host client on standalone ESXi will require that DHCP be available to the Brain appliance when booting for an IP to be assigned. After the system is licensed and the CLI of the Brain is available for login, if a Static IP is required, the initial DHCP setting for the management port can be changed to a static assignment using the CLI of the Brain.
 - Using vCenter with a vSphere client to do the deployment in VMware allows for either a Static or DHCP address assignment for the initial boot of the Brain appliance.
- ▼ Browse to the IP assigned to the Brain’s management interface to see the initial boot status messages. The

status message screen will update but a manual refresh is required to display any new information.

- When you are able to see the “System Setup and Provisioning” screen, enter proxy information if required for your environment and then enter the “License configuration” screen.
- ▼ Copy the licensing request code from the Vectra UI.
- ▼ Back on the Vectra customer portal, paste the license request code into the Licensing Request form in the “Enter Authorization Code” box available at <https://support.vectra.ai/vectra/additional-resources>.
- ▼ Copy the license once generated and paste it into the Vectra UI “Licensing Information” box.
- ▼ If your Brain deployment will not be online (connected to Vectra’s provisioner/updater system, this is only supported for Quadrant UX deployments), check the “Offline” box to enable an offline deployment. This means all licensing functions will be done manually offline.
 - If the “Offline box is not checked prior to hitting the “Save” button, they deployment will fail and will need to be started over with a fresh deployment of the OVF template.
 - Offline updates mode is automatically enabled when selecting the offline deployment mode.
 - VMware Brains deployed in offline mode can **never** be updated online.
- ▼ Click “Save” on the licensing configuration screen.
- ▼ After the license is validated, the file system will be decrypted, a performance test will be run, the Brain will reboot, and the Brain will reach out to the Vectra provisioning server and complete provisioning (if the deployment is in online mode). Finally, a success message will be presented with a button to redirect to the main UI login screen. Offline Brains follow a similar process but do not need to communicate with the provisioning server and can validate the license locally.
 - For Respond UX deployments, please see the deployment process in the [Vectra Respond UX Deployment Guide](#). You should not login and configure anything in the Quadrant UX (which would be available at this point) if you will be performing a Respond UX deployment.
- ▼ Initial login credentials for the UI are **admin / changethispassword**. Initial login credentials for the SSH access to the CLI are **vectra / changethispassword**.
 - You will be asked to change the password after the initial login.
- ▼ Complete your configuration using instructions available in the [Vectra Respond UX Deployment Guide](#) or [Vectra Quadrant UX Deployment Guide](#).

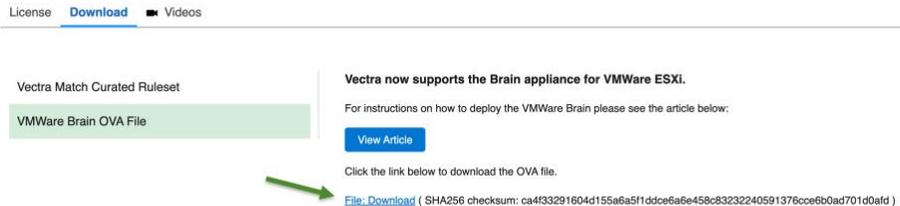
Brain Deployment in VMware

Requirements

- ▼ IP address and subnet mask for the Management interface of the Brain.
- ▼ DNS server addresses.
- ▼ Current login to a fully approved Vectra Support Portal account.
 - Accounts that are self-registered and not fully approved on the Vectra Support Portal will not have the license request option enabled.
- ▼ An open Proof of Value (Proof of Concept or Trial) that you are working with Vectra or a Vectra partner or a valid entitlement to Vectra NDR through purchase.
 - The licensing system cannot provide licenses for customers who are not currently entitled to a license through a trial or purchase.
- ▼ Per the [VMware Brain Requirements and Throughput](#) section, 32 core vSensors may need their NUMA settings adjusted before the initial power on. If deploying a 32 core vSensor, please see [32 Core NUMA Configuration](#) for details on checking and setting (if required) the NUMA autosize parameter.

Downloading the latest VMware Brain OVA image

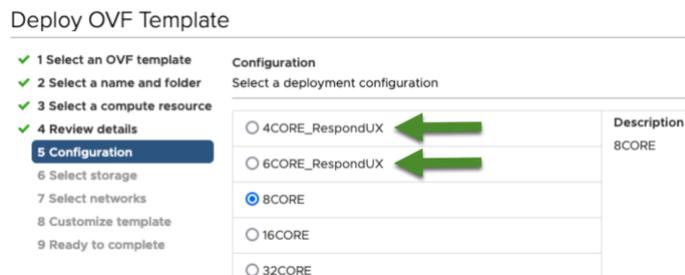
The current Brain OVA image can be downloaded from the Vectra Customer Support Portal after logging in. The URL for the download page is: <https://support.vectra.ai/vectra/additional-resources> and then click on the “Download” tab. Click on the “VMWare Brain OVA File” and then the “Download File” link to download the image. A SHA256 hash is also provided to allow you to verify the download completed successfully. Always download a current copy when you go to deploy a new VMware Brain for your organization. This will save time during deployment as fewer updates will need to be downloaded after deployment. Make this file available via a URL or the local filesystem where you will run the vSphere client from.



!! Choose one of the following two deployment methods for the OVA:

Deploying the OVA (vSphere Client / vCenter Server)

- ▼ On the host you wish to deploy the Brain on, right click and select “Deploy OVF Template...”
- ▼ Select the URL or Local file option depending on where you made the image available.
 - You can select the OVA itself or if you have chosen to decompress the OVA you can select the .ovf and associated .vmdk files.
 - Click “Next”.
- ▼ Configure a Virtual machine name and location for the virtual machine.
 - Click “Next”.
- ▼ Select a compute resource for the deployment.
 - Click “Next”.
- ▼ Review details and then click “Next”.
- ▼ Choose a Configuration and click “Next”.
 - Vectra may add additional configuration options in the future. Please refer to [VMware Brain Requirements and Throughput](#) for details on the supported configurations.
 - When deploying a v8.1 and higher base image, new 4 and 6 core Respond UX specific configurations are available. ONLY choose these if you are doing Respond UX for Network deployment.



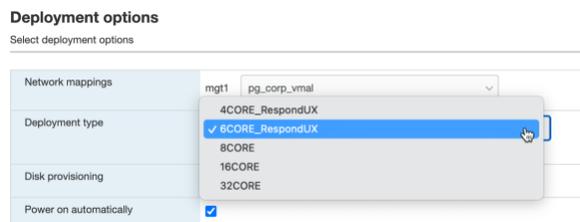
- ▼ Select storage.
 - Vectra recommends thick provisioning for storage (lazy or eager zeroed). Thin provisioning may work in some situations such as lab systems that don’t require high throughput, etc.
 - Storage DRS is not supported and should be disabled for this virtual machine.

- During the initial booting process, Vectra does a performance test to determine how well the system performs in comparison to established baselines.
 - Results can be retrieved from the CLI using the “`performance-test`” command when logged in as the “`vectra`” user. Additional detail is available in [Post Deployment Guidance](#).
 - ▼ Select the network for the mgt1 (Management) interface and click “Next”.
 - ▼ On the “Customize template” screen, please fill in the required details:
 - **DHCP** – Choose this box if you want the mgt1 (Management) interface to boot with DHCP enabled. If this option is chosen, the rest of the fields do not need to be filled in if DHCP will assign all of them.
 - **Hostname, IP Address, Netmask, Gateway, and DNS Servers** – Fill these in as required.
 - If a static IPv6 address is assigned during deployment, IPv6 support will be automatically enabled. Please see [IPv6 Management Support for Vectra Appliances](#) for more details.
 - **RespondUX** – Choose this option if you are doing a Respond UX for Network deployment.
 - When this option is selected, the Brain will boot directly into a state that is ready to be linked to the Vectra Cloud for use with the Respond UX. There will be no local Quadrant UX GUI served as there normally would be for a standard VMware Brain deployment before it is linked with Vectra for use with the Respond UX. Vectra personnel will still need to link your Brain to your Respond UX tenant to complete your deployment.
- RespondUX**

Select this to boot the VHE directly into RespondUX mode. This enables pairing with the Vectra AI Platform utilizing the Respond UX.
- This option should be selected for any Respond UX for Network deployment. i.e. You still need to pick this option even if you previously chose the “6CORE_RespondUX” configuration.
 - Click “Next”.
 - ▼ On the “Ready to complete” screen, validate all the details and when ready click “Finish”.
 - ▼ The OVF package will be imported and deployed.

Deploying the OVA (embedded host client for ESXi)

- ▼ Select “Create/Register VM” on your host. This will open a “New virtual machine” window.
- ▼ Select “Deploy a virtual machine from an OVF or OVA file and click “Next”.
- ▼ Enter a name for the Virtual machine name and then select or drag/drop your downloaded .ova file.
 - Click “Next”. If you see an error message about ignoring a disk, this can be ignored.
- ▼ Select the storage location for your VM and click “Next”.
- ▼ On the Deployment options screen, configure the following:
 - Network mappings – Choose the vSwitch to deploy the mgt1 (Management) interface into.
 - As mentioned in the [Licensing and Deployment Overview](#), DHCP is the only option supported by VMware when using the embedded host client for ESXi.
 - Deployment type – Choose the configuration you wish to deploy. Please refer to [VMware Brain Requirements and Throughput](#) for details on the supported configurations.
 - When deploying a v8.1 and higher base image, new 4 and 6 core Respond UX specific configurations are available. ONLY choose these if you are doing Respond UX for Network deployment.



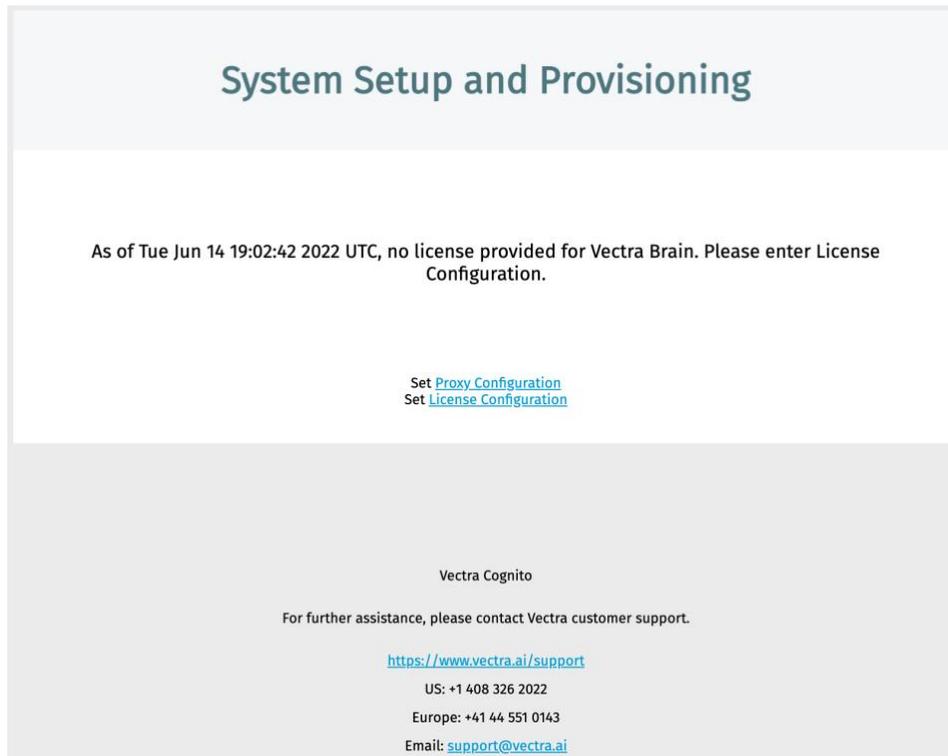
Initial boot up and licensing

After deploying the Brain VM into your VMware environment using one of the two methods above, it will need to be licensed during its initial boot process.

- ▼ Power on your Brain VM.
- ▼ Once the UI is available (this will be a few minutes after power on), use your browser to connect to your Brain VM (using the IP assigned statically, via DHCP, or via hostname if your Brain is in your DNS). If watching your console and you hit ESC when you see the Ubuntu boot screen you will see system messages. This particular message means that soon the UI will be available for license configuration:

```
[ *** ] A start job is running for Encrypted File Sys..ncrypted fs and send signal (41s / no limit)
```

- Navigate to `http://<your_brain_IP_or_hostname>` and you'll see a screen like this:



- ▼ At this point the initial boot process is paused until a valid license is entered. This process will not time out.
- ▼ If a proxy is required to communicate with Vectra for provisioning from your network, please enter the “Proxy Configuration” screen and enter your proxy information:
 - If this is not done now, it can be done after a license is saved but the provisioning process will time out and some time will be wasted until a proxy is configured.
 - **Please note:** This proxy configuration screen is only used to communicate with Vectra’s provisioning server and must utilize an HTTPS proxy. HTTP only proxies are not supported for this use. Other proxy configuration in the main Vectra UI (**Data Sources > Network > Brain Setup > Proxy & Status**) after deployment accepts HTTP proxies and is used by non-provisioning related items.
 - **Please note:** If you are doing a Respond UX deployment and require a proxy for non-provisioning related services and integrations (this includes linking to Vectra’s cloud for use with the Respond UX), you should configure that proxy at the CLI of your Brain AFTER you progress through this initial configuration and get to the “Success!” message at the end of this section. Please see the [Respond UX Deployment Guide](#) in the *Deployment > Proxy Support* section for more detail.

Proxy Configuration

Proxy Server Information

Enable Proxy

IP / Server name Port

Username

Password

Save Cancel

- ▼ Click into the “License Configuration” screen after saving a proxy configuration (if required):

License Configuration

License Request

eyJtb2RlbCI6IlZlRSIsInBsYXRmb3JlIjoidm13YXJlliw

Copy

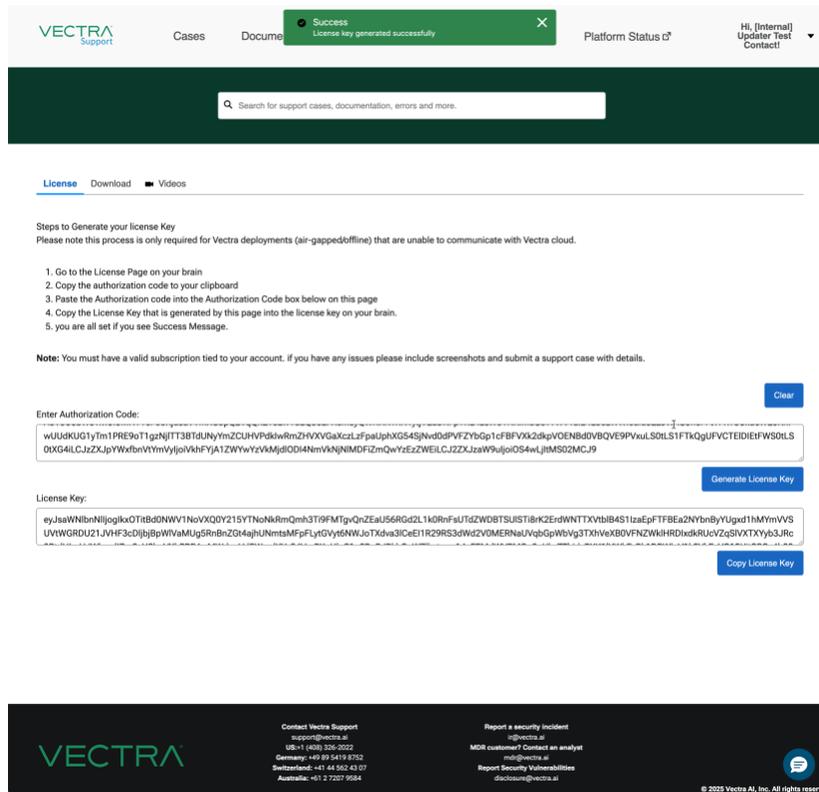
License Information

License

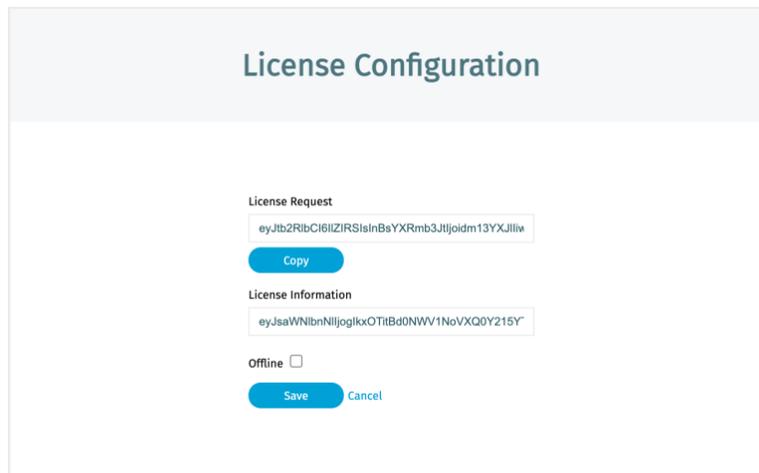
Offline

Save Cancel

- ▼ Copy the “License Request” information by clicking the “Copy” button.
- ▼ In another browser tab or window, navigate to <https://support.vectra.ai/vectra/additional-resources> and then the “License” tab.
 - If you are not already authenticated, you will be redirected to authenticate to your Vectra support account.
 - If you do not have a Vectra support account, you can self-register at the login screen, but licensing will not be available until your account is validated as being a Vectra customer or prospect involved in a trial.
- ▼ Paste the license information you copied into the “Enter Authorization Code” section of the page and click “Generate License Key”. You should get a “Success” message at the top and a key in the “License Key” box. Copy the license key using the “Copy License Key” button and go back to your Brain in your other tab or window.

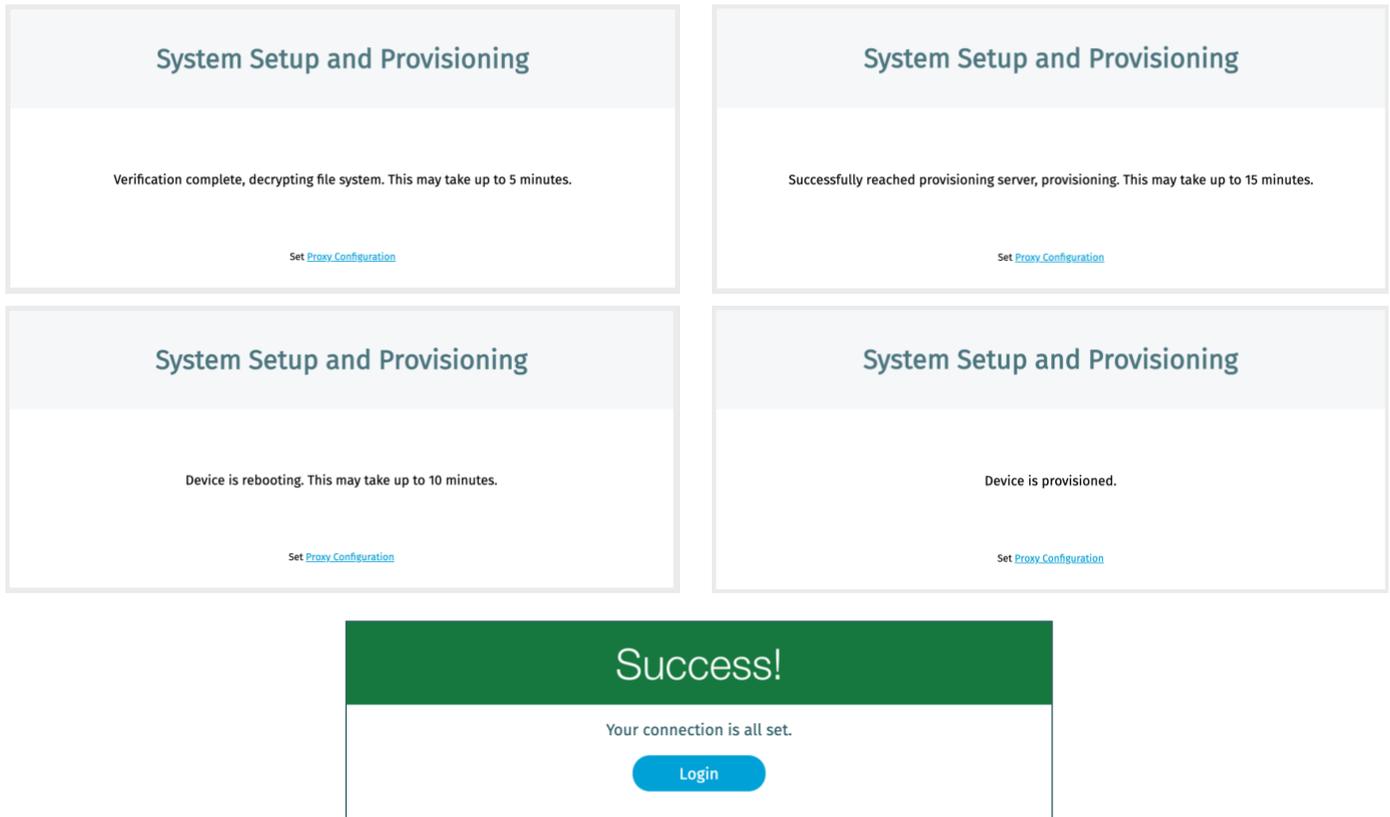


- ▼ Paste the license key into the “License Information” box.
 - **DO NOT CLICK SAVE YET!!!** You must first determine if your Brain will be “Online” or “Offline”.



- ▼ If you Brain will be “Offline” (Quadrant UX only) you **MUST** click the “Offline” checkbox now before clicking “Save” or your deployment will fail, and you will need to redeploy a new Brain VM into your VMware environment and start the process over.
 - Offline Brains do not communicate to Vectra’s provisioning service.
 - Offline Brains are typically used in “air gapped” environments where customers do not have internet access and the ability to communicate with Vectra directly.
 - For additional information regarding offline Brains see the following Vectra Support articles:
 - [Offline Updates \(v8.9+\)](#)
 - [Vectra Respond UX Deployment Guide](#) or [Vectra Quadrant UX Deployment Guide](#)
 - After determining if your Brain will be online (normal for the vast majority of customers) or offline, click “Save”. You may need to refresh the resulting page that may still say there is no license in place. After the page refresh, you should see screens like on the following page.

- ▼ After the license is validated, the file system will be decrypted, a performance test will be run, the Brain will reboot, and the Brain will reach out to the Vectra provisioning server and complete provisioning. Finally, a success message will be presented with a button to redirect to the main UI login screen. Offline Brains follow a similar process but will not need to communicate with the provisioning server and can validate the license locally.
 - Status messages will be updated but you must manually refresh the browser to see new messages.
 - Below are some examples:



- ▼ Once you see the success message and blue “Login” button you are ready to login to the main Vectra UI if performing a Quadrant UX based deployment.
 - Click the “Login” button and enter the default credentials.
 - Initial login credentials for the UI are: **admin / changethispassword**.
 - You will be asked to change the default password upon initial login.
- ▼ At this point, SSH to the CLI is available: Initial login credentials are: **vectra / changethispassword**
 - You will be asked to change the default password upon initial login.
- ▼ For Respond UX deployments, please follow the deployment process as detailed in the Vectra Respond UX Deployment Guide. You should NOT login to the Quadrant UX that is served from the Brain at this time!

Post Deployment Guidance

Setting a static IP and DNS after initial DHCP deployment

If you used DHCP for initial deployment but would like to configure a static IP for production use, you will need to login to the CLI of the Brain to set a static interface assignment. DNS for Brain VMs can be configured at the CLI or in the UI.

Logging in can be done via your hypervisor console function or using SSH to the management port if it was preconfigured with DHCP.

- ▼ Connect to your Brain CLI using your hypervisor console or "`ssh vectra@<IP or Hostname>`" if you use DHCP and already know the address or hostname.
- ▼ Once logged in to the Brain you can view command syntax for the "set interface" command:

```
set interface -h
Usage: set interface [OPTIONS] [mgt1] [dhcp|static] [IP] [SUBNET_MASK]
[GATEWAY_ADDRESS]

Sets mgt1 to either dhcp or static ip configuration

Options:
-h, --help Show this message and exit.
```

- ▼ Setting the IP address statically:
 - In v8.5 and higher of Vectra software, IPv6 is supported for the MGT1 interface. For full details, including information regarding dual stack support, please [IPv6 Management Support for Vectra Appliances](#) on the Vectra support portal. Below we will show how to enable IPv6 support (its off by default) and the syntax to use when setting an IPv4 or IPv6 address.
 - To enable/disable IPv6 support

```
# show ipv6 enabled
IPv6 is disabled

# set ipv6 enabled
Response: ok

# show ipv6 enabled
IPv6 is enabled

# set ipv6 disabled
Response: ok
```

- Setting IPv4 and IPv6 syntax examples:

Execute the following command to set the MGT1 interface to the desired static IP address:

```
IPv4 Syntax:
set interface mgt1 static x.x.x.x y.y.y.y z.z.z.z

Where:
x.x.x.x is the desired interface IP address
y.y.y.y is the desired interface network mask
z.z.z.z is the desired gateway

IPv6 Syntax:
set interface mgt1 static [IPv6 IP] [Subnet Mask] [Gateway]
```

```
Example:
set interface mgt1 static 2001:0db8:0:f101::25 64 2001:0db8:0:f101::1
```

- ▼ To change back to DHCP (default):

```
set interface mgt1 dhcp
```

- ▼ Configure DNS for the appliance:

Command syntax to set DNS (up to 3 nameservers are supported):

```
set dns [nameserver1 <ip>] [nameserver2 <ip>] [nameserver3 <ip>]
```

Example:

```
set dns 10.50.10.101 10.50.10.102
```

Verifying DNS Configuration:

```
show dns
```

- ▼ To set DNS in the UI, navigate to *Data Sources > Network > Brain Setup > DNS Entries* and edit the settings.

Setting static IP and DNS at the CLI Example:

```
vscli > set interface mgt1 static 172.16.12.11 255.255.255.0 172.16.12.1
Interfaces updated successfully
vscli > set dns 10.50.10.101
DNS Set: success
vscli > show interface
mgt1:
  Running:
    Gateway: 172.16.12.1,
    Ip: 172.16.12.11,
    Link Speed: 10Gbps,
    Link State: up,
    Mac: 00:0c:29:89:ad:a6,
    Mode: static,
    Netmask: 255.255.255.0
vscli > show dns
Id|Server      |Description
1 10.50.10.101 Configured DNS nameserver
```

Required Sensor update package for offline mode (Quadrant UX only)

As mentioned earlier in this guide, offline licensing and offline updates are configured separately. A Brain that is offline for licensing will not communicate with Vectra for updates so offline updates must be enabled. To be enabled for offline updates, you must enable manual updates for your deployment. This is typically done over a web meeting with screen sharing. Please see [Offline Updates \(v8.9+\)](#) on the Vectra support site for details on how to enable.

Once you are enabled for offline updates, you must apply a Sensor update package for virtual Sensors (cloud or customer hypervisor) to become fully functional. Any future updates will also require this package to have been installed. The [Offline Updates \(v8.9+\)](#) also provides details for how to perform offline updates.

Performance testing

As discussed earlier in this guide, a performance test is run during the initial boot process. This is to test the performance of the Brain against baselines that Vectra has established for the different configuration options.

Cached results from the initial performance test run can be retrieved from the command line while logged in as the “**vectra**” user. Additional performance tests can be run by using the `--force` switch on the performance test command. Please note the following:

- ▼ **Running the performance test is an intensive operation which takes down most services on the Brain.**
- ▼ Additional performance tests should only be run when your security team knows the Brain will be unavailable.
 - Paired Sensors will buffer metadata that can't be sent to the Brain so there should ultimately be no Detection gap, although this could introduce a delay in Detection publishing while the test is run.
- ▼ Baselines are set by Vectra for each of the various configurations of Brain.
 - Warning is for 10% below expectations. Critical is for 20% or more below expectations.
 - 260 MB/s is the minimum required throughput for all disks (OS and Data) and is represented by a score of 10.0 on the performance test in the “disk” category.
 - Critical is considered a failure and performance is not expected to be satisfactory. Vectra engineering considers systems which fail the performance test to be invalid configurations and customers should use more performant base hardware to ensure supportability, reliability, and performant operation.

Example:

```
vscli > performance-test --help
Usage: performance-test [OPTIONS]

Run a system performance test

Options:
--force Run all tests regardless of cached results.
-h, --help Show this message and exit.

vscli > performance-test
This may take up to five minutes. Most system services will be down for the duration of the test.
Test |Score |Result |Time
cpu 10.00 / 10.0 pass 30.04
cpu_steal 10.00 / 10.0 pass 0.06
disk 10.00 / 10.0 pass 47.94
memory 10.00 / 10.0 pass 0.00
memory_balloon 10.00 / 10.0 pass 0.05
overall 10.00 / 10.0 pass 78.09
```

Integrity checks

Vectra performs file system integrity checks to make sure that core libraries have not been altered. If the system detects changes during boot, a system setup and provisioning dialog will appear that is similar to the licensing screen.

- ▼ Click “Set File System Configuration”.
- ▼ Copy the “Error Code” and send it to Vectra support for decryption.
- ▼ Vectra has tooling to determine what has been changed, and if warranted can provide a whitelist code to the customer to allow the system to continue booting.
- ▼ Whitelist codes work one time. If the system again fails a file system integrity check, a new whitelist code will be required. Please work with Vectra support to ensure compliance.

Below are some example screenshots:



Configuration validation

During boot, the Brain determines which configuration it is running and sets some parameters differently depending on resource availability per configuration. This is an automatic process and requires no user input. Vectra may choose to support additional configuration options in the future. Please work with your Vectra account team to provide feedback regarding additional configuration options that would be useful to your organization. The `show system-health` command can be run at the command line as the `vectra` user to see that your configuration is a supported option. Look for the `[OK] VM Specifications`. The specific checks shown may not match your system. Vectra occasionally updates the specific checks used in the system-health command.

Example:

```
vscli > show system-health

===== Ran 8 check(s). 8 Passed, 0 Failed, 0 No Result =====

vscli > show system-health --verbose
[ OK ] Available Virtual Storage Space
[ OK ] Disk Writable
[ OK ] NIC Detection
[ OK ] Vectra User Password
[ OK ] Sensor Connectivity
[ OK ] Sensor Link Utilization
[ OK ] Sensor Tunnel
[ OK ] VM Specifications
```

License checks and renewal

Once a Brain is up and running, it will periodically check its license status. This will occur whether the Brain is online or offline (from the perspective of connection to Vectra). Once a Brain is 30 days from expiration, it will begin to send syslog messages with a count down until expiration. Once the license expires a new syslog message is sent (Quadrant UX). Respond UX deployments will write the message to the audit log which is available for query via API. Here are examples:

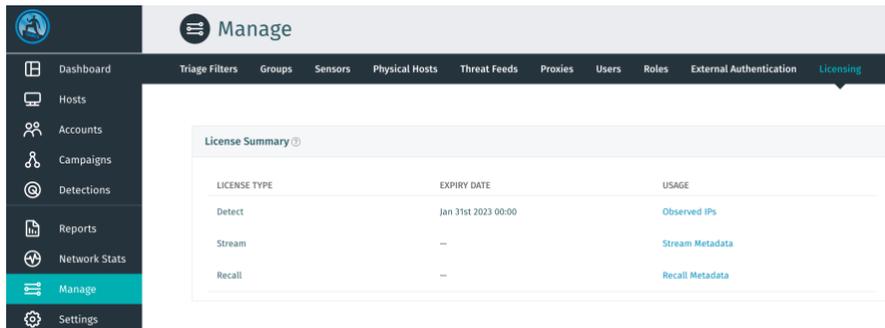
```
"License Checker: Detect License Expires in {days_until_expiration} days."
"License Checker: Detected Invalid/Expired License, disabling services"
```

The status of your license can be seen in the following locations in the Vectra UI:

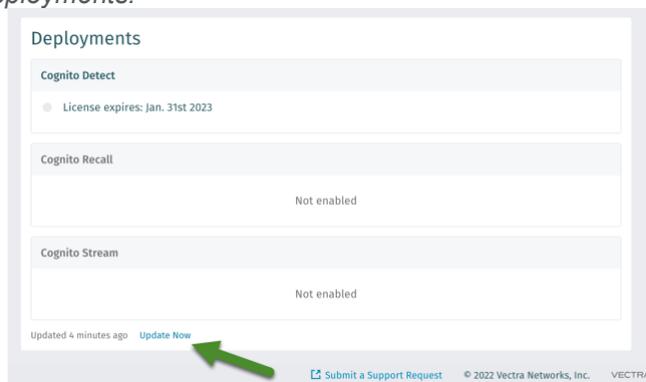
- ▼ *Manage > Licensing*
- ▼ *Discover > System Health > Deployments*
 - If your license status does not show, click the "Update Now" button at the bottom of the section.

Examples:

Manage > Licensing:

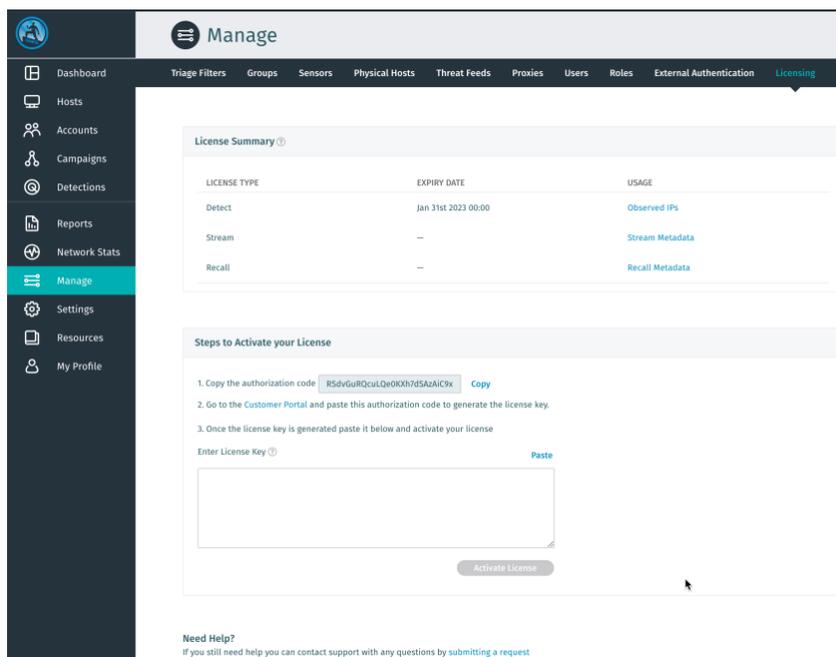


Discover > System Health > Deployments:



For Brains that are connected to Vectra, license renewal is an automated process that requires no user intervention. When your sales contract is renewed and the expiration date is updated, Vectra’s provisioning service will provide a new license key to your Brain.

If your Brain is offline (not connected to Vectra or air-gapped), to renew your Vectra license simply browse to the *Manage > Licensing* screen, copy the authorization code, provide it to Vectra (support, sales team, etc), and Vectra will provide you a new license key for entry into the UI once your entitlement is verified:



About VMware vCenter Integration

vCenter integration from the Vectra Brain for your deployment enables a number of features:

- ▼ “Virtual Infrastructure” view.
- ▼ vCenter host information artifacts help to feed Vectra’s automated Host ID.
- ▼ Additional VMware context is available for analysts on VMware hosts.
- ▼ vCenter alerts are possible as an additional notification type.

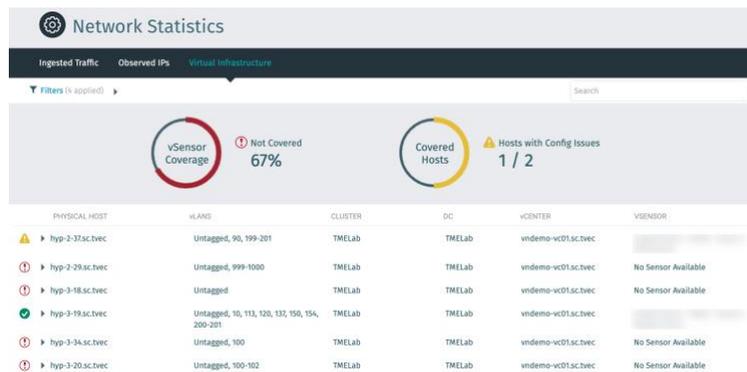
Virtual Infrastructure View

Enabling the vCenter API query connectivity helps with VMware vSensor deployment planning by identifying the physical hosts, clusters and data centers that currently have vSensor coverage, and those that do not have coverage.

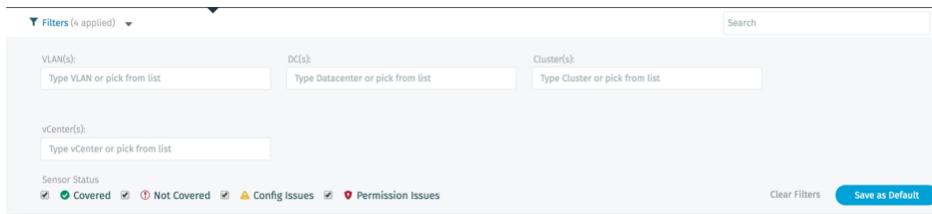
Enabling the vCenter connection also shows available resources on physical VMware hosts and exposes any configuration errors that might be affecting packet capture. This view, seen in your Vectra GUI at the *Network Stats > Virtual Infrastructure* page, helps the Vectra admin identify the exact requirements that need to be conveyed to VMware operational teams.

With this integration, the security team may not need to rely on the IT team to be notified of changes impacting them.

Once the vCenter integration is configured a *Network Stats > Virtual Infrastructure Hosts* page is enabled in the Vectra GUI:



The filter dropdown allows you to determine what is shown on the “Virtual Infrastructure” page:



A red exclamation point means that a particular physical hypervisor is NOT covered. This either means that there is no vSensor installed on the hypervisor or that the installed vSensor cannot be detected. A yellow warning sign icon means that there is a configuration issue with the installed vSensor:

PHYSICAL HOST	vLANs	CLUSTER	DC	vCENTER	vSENSOR
hyp-2-29.sc.tvec	Untagged, 999-1000	None	TMELab	vndemo-vc01.sc.tvec	vSensorCPG1-2-29b
BANDWIDTH 0.03 Mbps MEMORY Total: 256 GB Free: 236.1 GB CPU Total: 16 Cores Free: 15 Cores		DATSTORE hyp-2-29-ds02 SPACE Total: 3.3 TB Free: 2.7 TB DATSTORE hyp-2-29-ds01 SPACE Total: 924 GB Free: 923 GB VLANS Untagged, 999-1000 vSENSOR Missing VLANS			

A green checkmark means that the vSensor is configured and functioning properly:

✓ hyp-2-37.sc.tvec	Untagged, 90, 199-201	None	TMELab	vndemo-vc01.sc.tvec	vSensorCPG1-2-37b
BANDWIDTH 0.05 Mbps		DATASTORE ds02-2-37			
MEMORY Total: 128 GB Free: 101.1 GB		SPACE Total: 4.4 TB Free: 3.7 TB			
CPU Total: 16 Cores Free: 15 Cores		DATASTORE ds01-2-37			
		SPACE Total: 924 GB Free: 921.3 GB			
		VLANS Untagged, 90, 199-201			
		vSENSOR Paired			

vCenter Host ID artifacts

Vectra’s automated Host ID is a key benefit for analysts using the system. The goal of Vectra’s Host naming is to provide human-readable names associated with known Hosts.

Host names result from known information about the Host. Each observed name is referred to as an "artifact". Artifacts will typically be added to a Host record over time as more Host activity is seen and better associations are made. Host artifacts may be removed from a Host depending on the observed behaviors.

Hosts are tracked internally in a name agnostic manner. When assessing Host naming in your deployment, it is important to understand that Host names are decided at the time of viewing the web page. It is therefore expected that displayed Host names will change over time to reflect the most human readable name given the artifacts available at the time of page display.

The hostname obtained through vCenter/vSphere integration via an active query using the vCenter API is a key artifact when available in a customer environment. It is considered a best practice to enable the vCenter integration even if you will not deploy VMware vSensors in your environment.

For additional information regarding Vectra’s automated Host ID, please see the following Vectra support portal article:

[Understanding Vectra NDR Host Naming](#)

Additional Host context for analysts

When an analyst views a Host that is running VMware tools and reporting back to vCenter/vSphere with the Vectra vCenter integration enabled, additional context about the Host is available. To view this context, look at the left-hand side of the Host’s page in the Vectra UI for summary information including the VM name and operating system as reported by the vCenter API. The “**Host Details**” view has a more complete view. Below are examples of each:

VMWARE	
VMWare	
VM Name: Piper-desktop	Host Name: Piper-desktop
OS: microsoft windows 10 (64-bit)	UUID: 7d994ee6-562c-471f-b196-40ad89400760
	Physical Host: hyp-2-37.sc.tvec
	Cluster: TMELab
	Data Center: TMELab
	vCenter: vndemo-vc01.sc.tvec
	OS: microsoft windows 10 (64-bit)
	Power State: poweredOn
	Nets: IPS MAC
	192.168.150.100 00:50:56:93:b0:89

vCenter alerts

Once the vCenter integration is configured, additional alerts are available that are specific to changes in the environment that may merit security consideration. To enable these alerts navigate to *Settings > Notifications > Alert Emails*, select the “**Edit**” or pencil icon, scroll to the bottom of the Alert Emails settings, and enable the toggle to “**Send vCenter alerts**”. Some example scenarios where an alert will be sent are:

- ▼ A new physical hypervisor where a vSensor may be needed has been added to the environment.
- ▼ A vSensor has been moved or powered down.
- ▼ A VM is moved from a Host that is monitored by a vSensor to a Host that is not monitored by a vSensor.

Enabling vCenter Integration

Prepare vSphere account for Brain access

To connect the Brain to vSphere, a vSphere user account and password must be configured into the Brain. The vSphere user account must have at least global, read-only rights. The Brain will not attempt to write any data to your VMware environment.

To ensure that the vSphere user/group the Brain will use has global, read-only access, use the following steps in the vSphere UI:

- ▼ From the vSphere Administration page select *Access > Global Permissions*.
- ▼ Click the **plus** sign to display the global permissions dialog.
- ▼ At the bottom of the left pane, click **Add**.
- ▼ Ensure the domain is set to the proper domain, select the **users** or **groups** you intend to use in Vectra’s configuration to connect to vCenter’s API and click “OK”.
- ▼ In the **Assign Role** section, select **Read-Only** from the drop-down list.
- ▼ Make sure the **Propagate to children** checkbox is selected, and click “OK”.

Configure vCenter/vSphere integration

You will need to have the IP or hostname of your VMware vCenter server. You can configure multiple integrations if you have more than one server to connect to. You will also need to have the port number, username and password.

Navigate to *Settings > External Connectors > vCenter* and edit the vCenter settings. Any previously configured vCenters will be shown in this area:

HOST	PORT	USERNAME	PASSWORD
vndemo-vc01.sc.tvec	443	vectraapi

Click on the + Add vCenter to add an additional vCenter, fill in the blanks, and click “**Save**”

Resizing the Brain

In some environments, you may wish to start with a smaller Brain instance and then later move to a larger Brain instance to handle additional load (metadata coming from paired sensors or additional paired sensors).

- ▼ Please see: [Resizing Virtual Sensors and Brains](#) for details.

Next Steps

At this point your VMware Brain is fully deployed and you can move on to other tasks associated with your overall deployment.

It is recommended to follow the [Vectra Respond UX Deployment Guide](#) or [Vectra Quadrant UX Deployment Guide](#) for additional information regarding initial settings for your deployment. You may wish to deploy and pair network Sensors or configure other Vectra offerings such as Recall, Stream, CDR for M365, IDR for Azure AD, CDR for AWS, etc. Additional documentation can be found in the [Vectra Product Documentation Index](#) on the [Vectra Support](#) site.

Worldwide Support Contact Information

- ▼ Support portal: <https://support.vectra.ai/>
- ▼ Email: support@vectra.ai (preferred contact method)
- ▼ Additional information: <https://www.vectra.ai/support>